

MEMORY CONTROLLER

TECHNICAL FIELD OF THE INVENTION

5 This invention relates to a memory controller, and in particular to a controller for a SDRAM (Synchronous Dynamic Random Access Memory) device, although the invention is also applicable to other types of memory, and to a method of operation of a memory controller.

10

BACKGROUND OF THE INVENTION

Computer systems must be provided with sufficient data storage capacity to operate correctly. This data storage capacity is typically provided as Random Access Memory (RAM), and SDRAM is a common form of RAM.

However, the rate at which data can in practice be transferred from a SDRAM remains lower than the rate at which data can in theory be transferred. That is, each access request sent to a SDRAM memory chip, relating to a read operation, incurs a read latency.

When a master device makes multiple read access requests, this read latency can be incurred for each access request.

Accesses to the SDRAM chip are performed by a SDRAM controller, which typically takes the form of an integrated circuit which is separate from the SDRAM. The SDRAM controller is connected to the SDRAM by means of a memory data bus, and the SDRAM controller must operate as far as possible to maximize efficient use of the bandwidth of that bus.

SUMMARY OF THE INVENTION

It is an object of the present invention to minimise the time required to return all of the data read from the memory to the requesting master, while also ensuring efficient use of the bandwidth of the memory data bus.

5

More specifically, according to a first aspect of the present invention, a SDRAM controller determines for each received access request whether the required data can be retrieved in a single burst, or whether multiple bursts are
10 required.

The SDRAM controller forms a queue of bus access requests, and, if multiple bursts are required for a single read access request, the bus access requests relating to the
15 multiple bursts are queued effectively simultaneously, or at least on successive clock cycles.

This has the advantage that the overall performance of the computer system is optimized since a higher bandwidth can
20 be achieved on the memory data bus, thereby allowing the memory to be used more efficiently.

BRIEF DESCRIPTION OF DRAWINGS

25 Figure 1 is a block schematic diagram of a computer system in accordance with the present invention.

Figure 2 is a block schematic diagram of a SDRAM controller in the computer system of Figure 1.

30

Figure 3 is a flow chart illustrating a method in accordance with the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

35

Figure 1 is a block schematic diagram of a computer system 10. The general form of the system 10 is conventional, and will be described herein only to the extent necessary for a complete understanding of the present invention.

5

In the illustrated embodiment, the system 10 includes an application-specific integrated circuit (ASIC) 20, which includes various modules 25, such as a processor core (CPU) 27. These modules are interconnected by a bus 30, which
10 may advantageously be an AHB bus, but which can be any convenient form of bus.

However, the invention is not limited to such a structure. The invention is also applicable to a device such as a
15 programmable logic device (PLD) or field programmable gate array (FPGA), which can then be configured to contain multiple modules which act as bus masters. The device may then, but need not, contain an embedded processor.

20 Connected to the ASIC 20 is a memory chip 40, in the form of a Synchronous Dynamic Random Access Memory (SDRAM).

Accesses to the SDRAM 40 from the ASIC 20 are performed by a specific SDRAM controller 50 connected to the bus 30 in
25 the ASIC 20.

Again, the invention is not limited to such a structure. The SDRAM controller 50 may be integrated with the bus masters in a single device, or may be provided as a
30 separate device.

The SDRAM controller 50 is connected to the SDRAM 40 by way of a memory bus 60, which in the illustrated embodiment of the invention is also an AHB bus.

35

Figure 2 is a block schematic diagram, showing the form of the SDRAM controller 50.

5 The SDRAM controller 50 is shown in Figure 2, and described herein, only to the extent required for an understanding of the present invention. Other features of the SDRAM controller, which are not required for that purpose, will not be described, and may be assumed to be generally conventional, as known to the person of ordinary skill in
10 the art.

In the illustrated embodiment, the SDRAM controller 50 has multiple bus interface blocks 52, for connection to respective bus master devices. For example, in the system
15 shown in Figure 1, there may be one bus interface 52 allocated for connection to each of the modules 25 and the CPU 27. However, in other embodiments of the invention, there may be only one such bus interface block.

20 Memory access requests, received by the SDRAM controller 50 at the bus interface blocks 52, are passed to a control logic block 54, the operation of which is described more fully below.

25 After processing in the control logic block 54, the memory access requests are placed in a queue in a queue store block 56, which may for example take the form of a first-in, first-out memory. The memory access requests from the queue are then passed in turn to a SDRAM interface block
30 58.

Figure 3 is a flow chart, illustrating a method performed in the control logic block 54, according to an aspect of the present invention.

The process starts at step 300, when a read access request is received at a bus interface 52 from one of the master devices.

5 The read access request indicates the amount of required data with reference to the properties of the AHB bus 60, namely the burst length, which is a feature of the bus protocol, and the AHB word size, which can be less than or equal to the width of the bus. The read access request
10 also indicates the burst type, i.e. whether a wrapping burst or an incrementing burst is required.

Also in step 300, the control logic 56 reads the starting address of the request, that is, the address within the
15 SDRAM 40 from which data is first to be retrieved.

In step 308, the control logic 56 then calculates the number of SDRAM bursts required to fulfil the access request.

20 For example, if the AHB word size is 64 bits, and the AHB burst length is 16, while the SDRAM word size is 32 bits, and the SDRAM burst length is 8, then four SDRAM bursts are required to fulfil the access request if the starting
25 address of the request corresponds with a SDRAM burst boundary, while four or five SDRAM bursts are required, depending on whether the burst type is wrapping or incrementing, if the starting address of the request does not correspond with a SDRAM burst boundary.

30 As another example, again taking the SDRAM word size to be 32 bits, and the SDRAM burst length to be 8, if the AHB word size is 32 bits, and the AHB burst length is 8, but the starting address of the request does not correspond
35 with a SDRAM burst boundary, then two SDRAM bursts are required to fulfil the access request if the access request

indicates that the required burst type is incrementing, as opposed to wrapping.

Next, the control logic 54 determines the starting SDRAM
5 addresses of the required SDRAM bursts. Thus, in step 312,
the control logic then translates the AHB address into a
SDRAM address, within the SDRAM 40. In the case discussed
above, where more than one SDRAM burst is required, the
control logic calculates the SDRAM address for the first
10 SDRAM burst. The SDRAM address is made up of a SDRAM chip
select, a SDRAM row address and a SDRAM column address.

Thus, a set of SDRAM devices returning data for a
particular request comprise a physical bank of memory.
15 Multiple physical banks may be provided, in which case each
physical bank is accessed using a different chip select.
Within the physical bank, a specific memory location is
defined by a row address and a column address. The
calculated SDRAM address therefore uniquely identifies a
20 memory location within the memory device.

In step 314, in the case where more than one SDRAM burst is
required, the arbiter 100 also determines SDRAM addresses
for the remaining SDRAM bursts.

25 In each case, since the chip select and starting row
address values remain the same throughout an AHB burst, the
calculated starting SDRAM addresses for the second and
subsequent SDRAM bursts need relate only to the column
30 address values.

The separate read requests for each required SDRAM burst,
including the respective starting addresses, are then
placed into a queue of access requests in the queue store
35 56 of the SDRAM controller 50. The stored access requests
are then handled in turn by the SDRAM interface 58.

As is known to the person skilled in the art, the control logic 56 may also, in addition to the processes described herein, apply a form of prioritisation to the access
5 requests when placing them into the queue of access requests in the queue store 56. For example, access requests received on different bus interfaces 52 may be given different priorities.

- 10 In addition, or alternatively, access requests received on different bus interfaces 52 may be prioritised in a way which maximises the efficiency of use of the memory bus 60. For example, opening a page of the SDRAM to process an access request results in a delay in processing.
15 Therefore, it is advantageous if access requests relating to the same page of the SDRAM can be queued consecutively.

In accordance with the invention, therefore, the SDRAM controller can ensure that, for an AHB burst corresponding
20 to multiple SDRAM bursts, the access requests corresponding to the multiple SDRAM bursts are queued together, so that the read latency is incurred only once. When placing the multiple access requests in the queue store 56, a flag may be set on at least the first of said stored access
25 requests. This allows the SDRAM controller 50 to perform back-to-back SDRAM read bursts, and therefore increases the effective rate at which data can be read from the SDRAM.

Further, when enough SDRAM read bursts have been performed
30 to retrieve all of the data requested in the access request received on the bus interfaces 52, no additional data need be retrieved.

The invention has been described herein with reference to
35 one particular embodiment. However, other embodiments of the invention are also possible. The scope of the present

invention is therefore to be determined only by the accompanying claims.